

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-03-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Services, Directorate for Information Operations and Reports (DOD Form 4302). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not have a valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

0232

maintaining the
reducing
22202-
a currently

1. REPORT DATE (DD-MM-YYYY) 15-07-2001		2. REPORT TYPE Final		3. DATES COVERED (From - To) 15-09-2000 to 15-06-2001	
4. TITLE AND SUBTITLE Biologically Inspired Direct Adaptive Guidance and Control for High-Bandwidth Flight Systems				5a. CONTRACT NUMBER F49620-00-C-0055	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Eric Corban, Cole Gilbert, Anthony Calise, Allen Tannenbaum				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Guided Systems Technologies, Inc. P.O. Box 1453 McDonough, GA 30253-1453				8. PERFORMING ORGANIZATION REPORT NUMBER GST-00-3.1	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research 801 N. Randolph Street Arlington, VA 22201-1977				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approve for Public Release: Distribution Unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The capability of biological flight systems to autonomously maneuver, track, and pursue evasive targets in a cluttered environment is vastly superior to current engineered systems. The reported effort seeks to fully characterize the tracking, guidance and control functions of one such biological system (the male flesh fly in pursuit of the female), and to use the developed understanding to improve the design of engineered guidance and control systems for small autonomous air vehicles. The phase I effort produced demonstration of feasibility in three distinct areas: development and demonstration in simulation of a direct neural network adaptive guidance law for intercept; development and demonstration of experimental methods to capture and quantify the trajectories and orientation of the insects in pursuit/evasion scenarios; and further development and demonstration of practical image processing algorithms for implementation of the developed guidance system concepts. The proposed phase II program will extend and expand the work in these three areas to produce practical biologically-					
15. SUBJECT TERMS Adaptive Guidance and Control, Image Processing, Flesh Fly Pursuit and Evasion					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

20030709 015

Biologically Inspired Direct Adaptive Guidance and Control for High-Bandwidth Flight Systems

**Final Report
July 15, 2001**

Report Number GST-00-3.1

Reporting Period 9/15/00- 6-15-01

SBIR Phase I
Contract F49620-00-C-0055

Prepared for:

Robert H. Cohn
AFOSR/NL
801 N. Randolph Street, Room 732
Arlington, VA 22201-1977
Phone: (703) 696-7722
email: Robert.cohn@afosr.af.mil

Prepared by:

J. Eric Corban, Cole Gilbert, Anthony J. Calise, Allen R. Tannenbaum

Guided Systems Technologies, Inc.

P.O. Box 1453
McDonough, Georgia 30253-1453
corban@mindspring.com
(770) 898-9100

PREFACE AND ACKNOWLEDGEMENTS

The work presented herein was funded by the US Air Force Office of Scientific Research under STTR Phase I Contract F49620-00-C-0055 and coordinated with AFRL/MNG at Eglin AFB. The phase I proposal was submitted in response to the 2000 Air Force STTR solicitation topic A00T012 entitled "Implementation of Biomimetic Precision Flight in Autonomous Vehicles." The work reported herein depends in part on previous research in many diverse areas, and the contributions of the team members and their previous research sponsors in each of these areas is gratefully acknowledged. The phase I experimental work also benefited greatly from the commercial products being developed by Dr. P.O. Zanen of Synceros, Inc. in association with the Cornell Laboratory. This support is greatly appreciated. The contributions of research assistant Randy Weinstein, who supported the program while a student at Georgia Tech, are also appreciated.

TABLE OF CONTENTS

REPORT DOCUMENTATION PAGE	i
TITLE PAGE	ii
PREFACE AND ACKNOWLEDGEMENTS	iii
TABLE OF CONTENTS	iv
1. INTRODUCTION	1
1.1 PROBLEM AND OPPORTUNITY	1
1.2 ORGANIZATION OF THE REPORT	2
2. BACKGROUND ON HIGH BANDWIDTH ADAPTIVE CONTROL	3
2.1 BACKGROUND ON NEURAL NETWORK AUGMENTED CONTROL	3
2.2 REVIEW OF NEURAL NETWORK ADAPTIVE CONTROL	3
3. ADAPTIVE GUIDANCE LAW DEVELOPMENT	9
3.1 REVIEW OF PROPORTIONAL NAVIGATION	9
3.2 REQUIRED ACCELERATION	10
3.3 THE EFFECT OF TARGET MANEUVERS	12
3.4 NEURAL NETWORK BASED ADAPTIVE GUIDANCE	13
3.5 NUMERICAL RESULTS	14
4. CHARACTERIZATION OF FLESH FLY PURSUITS	19
5. IMAGE PROCESSING ALGORITHMS	22
6. SUMMARY, CONCLUSIONS, AND RECOMMENDATIONS	29
REFERENCES	30

1. INTRODUCTION

1.1 PROBLEM AND OPPORTUNITY

Rapid advances in many technology areas will soon make possible small or micro air vehicles that can be deployed in large numbers at low cost. In theory, military and commercial applications of such vehicles abound. However, in practice these vehicles must achieve a high level of autonomy in operations to be successfully deployed in large numbers. Furthermore, much of their utility is to be derived from operations in densely cluttered environments that include unknown obstacles. Traditional guidance and control technology is inadequate to meet the demand for autonomy in such environments, and dependence on traditional design methods will lead to expensive solutions that demand tremendous computational resources. There are, however, a multitude of biological systems that, with relatively crude sensors and limited computing resources, operate very successfully in such environments. This program seeks to emulate the success of these biological systems in an engineered guidance and control system. This is to be accomplished by developing a complete understanding of the tracking, guidance and control functions of a candidate biological system (the flesh fly), and to then use this understanding to inspire, develop and demonstrate effective and affordable guidance and control technologies for autonomous operation of small or micro air vehicles.

The STTR program solicitation issued by the Air Force calls for a demonstration of "the use or incorporation of insect neurobiology or sensory processing into the design of precision flight capability of autonomous air vehicles." More specifically, the solicitation calls for development and demonstration of biologically inspired guidance, navigation and control (GNC) sensors, components, and/or systems that will enable an autonomous small air vehicle to search for, detect, pursue, and rendezvous with an evasive target in a densely cluttered environment.

As stated in the solicitation, desirable characteristics of the GNC system include:

- Tolerance to transient sensor information distortion or obscuration;
- Capable of obstacle avoidance (e.g. buildings or trees, power lines, or overhanging branches);
- Detection, acquisition, tracking and guidance to a moving target in background clutter;
- Precision rendezvous with the target (e.g. warhead event, electronic tagging);
- Flight control robustness to wind gusts and turbulence near structures.

Guided Systems Technologies, Inc. (GST) assembled a uniquely qualified team of researchers to address these requirements. The phase I program was designed to demonstrate the feasibility and merit of combining university research in the fields of entomology, target tracking, aerospace vehicle guidance and control, as well as image processing and computer vision, with the practical systems integration and UAV capabilities to meet the challenge. The phase I

program produced significant results in three distinct areas that are to be extended and combined in phase II to produce a demonstration of the solicited capability.

1.2 ORGANIZATION OF THE REPORT

Chapter 2 presents a summary of neural network adaptive output feedback control and illustrates its use in design of a high bandwidth flight control system for unmanned helicopters.

Chapter 3 presents a novel method to extend the existing neural network adaptive control methods into the realm of intercept guidance law development. The chapter begins with a review of the classic proportional navigation guidance law. An adaptive form of the guidance law is then developed. The chapter concludes with the presentation of simulation results for intercept of a target.

Chapter 4 documents the phase I work completed with regard to experimental characterization of flesh fly pursuit and evasion. The chapter includes a description of the experimental methods and facilities, as well as presentation of results obtained to date.

Chapter 5 presents a compilation of work performed in the area of image processing.

Chapter 6 completes the report with a summary of the phase I effort, draws conclusions from the work completed, and makes recommendations for continued research.

2. BACKGROUND ON HIGH BANDWIDTH ADAPTIVE CONTROL

2.1 Background on Neural Network Augmented Control.

Researchers at the Georgia Institute of Technology and GST have, over the last decade, developed very effective methods for realizing practical neural network adaptive control systems for flight vehicles. For example, in the USAF RESTORE (reconfigurable systems for tailless fighter aircraft) program, an adaptive command augmentation system based on full state feedback was used to stabilize an airframe and maintain acceptable handling qualities in a wide variety of simulated control surface failures, and was successfully flight tested on the X-36 unmanned test aircraft. This same adaptive control technology has been applied to autopilot design for precision guided munitions, and was flight demonstrated on the MK-84 JDAM in 2001. GST and Georgia Tech also have an active program to flight test a recently developed output feedback formulation of the neural network adaptive law on a small agile unmanned helicopter. The output feedback result, when combined with a technique known as pseudo control hedging, allows for the design of a very-high bandwidth controller that can be used to exploit the full operational envelop of the unmanned flight vehicle in a manner similar to that achieved by a skilled human pilot. This neural network adaptive control system technology is a key element of the foundation on which the subject work of adaptive guidance is based, and is therefore reviewed in some detail in this chapter.

2.2 Review of Neural Network Adaptive Control

Both classical and modern control design methods are fundamentally limited by the presence of unmodeled high frequency effects, and control nonlinearities due to saturation. The same is true in adaptive methods that do not attempt to learn and interact with these effects. Yet biological systems routinely exhibit high bandwidth performance in the presence of these effects.

Therefore, one goal of the proposed research in control is to demonstrate an approach to flight control design that is inspired by the performance observed in nature. The approach is illustrated below using results for the design of pitch-angle flight control system of an experimental unmanned helicopter.

This section presents a novel approach to adaptive output feedback control¹. Of particular interest here is its use in high bandwidth flight control, which is made possible through interaction with (as opposed to avoidance of) poorly modeled high frequency dynamics, and saturation. Adaptation is achieved using only input/output sequences of the uncertain system. Furthermore, adaptation continues under saturation (position and/or rate) and the methodology permits other strong nonlinearities that may due to mixed utilization of continuous and discrete (bang-bang) control.

Neural Network-Based Adaptive Output Feedback Control: Consider the dynamics of an observable nonlinear plant represented by the following non-affine SISO system

$$\dot{x} = f(x, u), y = h(x) \quad (1)$$

where $x \in \mathbf{R}^n$, $u, y \in \mathbf{R}$ and $f, h \in \mathbf{C}^\infty$ may be partially known or unknown functions. Assume that n and the relative degree(r) are known. With the knowledge of n and r , we can postulate an equivalent expression for the system in a normal form as

$$\begin{aligned} y^{(r)} &= h_r(\phi, \eta, u) \\ \dot{\eta} &= f_0(\phi, \eta) \end{aligned} \quad (2)$$

where $\phi = (y, \dot{y}, \dots, y^{(r-1)})^T \in \mathbf{R}^r$, and $\eta \in \mathbf{R}^{n-r}$ is the state vector associated with the zero dynamics

$$\dot{\eta} = f_0(0, \eta) \quad (3)$$

which are assumed to be stable. It is often the case that an approximate expression exists for $h_r(\cdot)$ in Eq.(2) which can be formulated using only the instantaneous available measurements. This approximation can be used as a basis for an inverting control law by defining

$$\hat{h}_r(y, u) = v \quad (4)$$

where v is the so-called pseudo control variable. It is assumed that this function can be inverted to obtain the actual control in the form

$$u = \hat{h}_r^{-1}(y, v) \quad (5)$$

Applying the approximate inverse into Eq.(2), we obtain

$$y^{(r)} = v + \Delta \quad (6)$$

where $\Delta = h_r(\phi, u) - \hat{h}_r(y, u)$.

Figure 2.1 depicts the structure of this closed-loop neural network-based adaptive control system. The adaptive control term, v_{ad} , is introduced to cancel the model inversion error Δ . From the universal approximation property of multilayer NNs, a single-hidden layer NN is selected to approximate Δ . The structure of the NN, presented in Figure 2.2, is expressed as $W^T \sigma(V^T x)$ where W, V are matrices containing the interconnection weights, and $\sigma(\cdot)$, shown in Figure 2.3, is an activation function of the hidden nodes.

The following weight update rules in Equation (7) guarantee the uniform ultimate boundedness of the error dynamics and NN weights.

$$\begin{aligned}\dot{\hat{V}} &= -\Gamma_v \left[\eta (z \hat{W}^T \hat{\sigma}') + \lambda_v \hat{V} \right] \\ \dot{\hat{W}} &= -\Gamma_w \left[(\hat{\sigma} - \hat{\sigma}' \hat{V}^T \eta) z + \lambda_w \hat{W} \right]\end{aligned}\quad (7)$$

Control Hedging: In order to exploit the full maneuvering capability of a flight vehicle with adaptive control system design, it is necessary to protect the adaptive process from adapting to characteristics such as control saturation. The approach employed here is referred to as Pseudo-control Hedging. In this approach, the adaptation is protected from the effects of hard nonlinearities by shifting the dynamics in the command filter portion of the control system architecture backwards by an estimate of the amount that the real system is unable to achieve due to hard nonlinearities.

This approach can be used to:

- Isolate the effects of hard nonlinearities from the adaptation
- Mitigate the effects of approximately known input dynamics (such as actuator dynamics) from the adaptation
- Allow the adaptation to handle a class of unknown hard nonlinearities (such as an actuator failure).

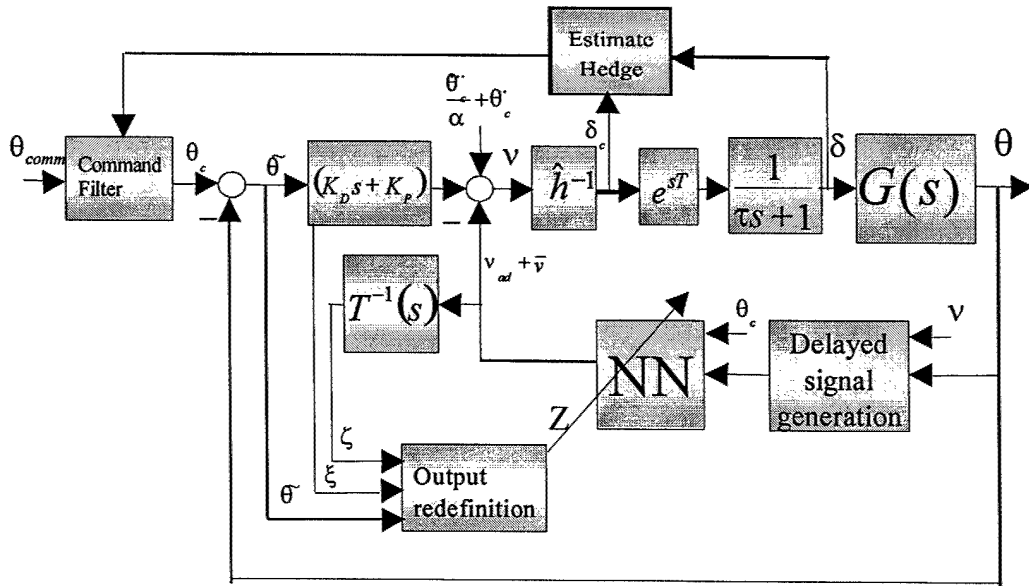


Figure 2.1 Block diagram of neural network adaptive controller – the output feedback formulation without an observer.





Figure 2.2 Topology of a single-hidden layer neural network and a typical activation function.



$$\sigma(z) = \frac{1}{1 + e^{-az}}$$

Figure 2.3 Typical activation function employed in each node of the hidden layer, plotted for different values of the internal activation potential, a .

Design Application to Pitch-attitude Control of an Unmanned Air Vehicle: Consider the design of a flight control system for a small unmanned helicopter. In the practical design of such a vehicle it is common to include a small control rotor. In this case, control inputs to the main rotor are instead fed to the control rotor, and the main rotor control inputs are a result of the control rotors reaction to both command and flight conditions. The control rotor is thus a mechanical feedback device and can be used by the designer to impart desired characteristics on the response of the overall dynamic system. While these same characteristics can be imparted through feedback with an electronic control system, it is desirable to retain the control rotor for manual flight operations such as may be required in response to an electronic system failure or software fault on an experimental control system testbed. The control input to the main rotor is affected by the control rotor flapping angle, actuator dynamics and time delay which add additional dynamics within the desired bandwidth of the control system design. The pitch channel equation of motion of such a helicopter can be expressed as a following SISO linearized system:

$$\begin{bmatrix} \dot{u} \\ \dot{q} \\ \dot{\theta} \\ \dot{\beta} \end{bmatrix} = \begin{bmatrix} X_u & X_q & X_\theta & X_\beta \\ M_u & M_q & 0 & M_\beta \\ 0 & 0.9991 & 0 & 0 \\ Z_u & -1 & 0 & Z_\beta \end{bmatrix} \begin{bmatrix} u \\ q \\ \theta \\ \beta \end{bmatrix} + \begin{bmatrix} X_\delta \\ M_\delta \\ 0 \\ Z_\delta \end{bmatrix} \delta \quad (8)$$

$y = \theta$

where y is the controlled variable. The state vector $x = (u, q, \theta, \beta)^T$ where u is forward velocity, q is pitch rate, θ is pitch angle, β is control rotor flapping angle and the control variable δ is the longitudinal cyclic input. The bandwidth limitation is largely due to control rotor dynamics (at approximately 8-10 rad/s), and filtering and digital processing time delays (between 0.02 to 0.04 seconds) that are present in the control path from the rate sensor to the main rotor swash plate. Moreover, there are significant position and rate limits in the actuator. The approximate inversion law

$$\delta_c = \hat{h}^{-1}(q, v) = \frac{\tau}{\hat{M}_\delta} (\alpha v - \hat{M}_q (\hat{M}_q + \alpha) q) \quad (9)$$

corresponds to having the desired open loop transfer function $G_d(s) = \frac{y(s)}{v(s)} = \frac{\alpha}{s^2(s + \alpha)}$.

The control deflection is assumed to respond to the input command with 0.03 second time delay and actuator dynamics where its time constant is selected as 0.04 sec. The control design locates the poles of the error dynamics at $-8 \pm 6i$, -20 . This is well within the bandwidth of the flapping dynamics, control rotor dynamics, and at a frequency that results in significant phase shift due to time delay.

Figure 2.4 compares the pitch attitude tracking performance with and without the neural network.

The tracking performance without adaptive neural network results in a limit cycle of the tracking error. On the other hand, excellent tracking performance is achieved using the adaptive neural network. Moreover, the rate limit is encountered whenever there is a reversal in the command. This result exhibits the fact that the adaptive controller achieves improved performance by interacting with the control rotor dynamics, actuator dynamics, time delay and control limits. This interaction is a result of a highly nonlinear process, and is difficult to achieve using a linear controller.




Figure 2.4 Pitch attitude command and response time histories with and without neural network augmentation.

3. ADAPTIVE GUIDANCE LAW DEVELOPMENT

A primary phase I objective was to develop a direct adaptive approach for correcting a proportional navigation (PN) guidance law for uncertainty associated with the target acceleration.

The term “direct” is used in this context to identify the fact that the developed approach does not depend on the traditional use of an observer. Observer-based solutions rely on a model for target behavior, whereas a direct adaptive approach that is neural network (NN) based can be free of target modeling assumptions, relying instead on the universal approximation properties of artificial neural networks to model the target’s evasive strategy. Such an approach is expected to be less vulnerable to intelligent targets that can take advantage of the modeling assumptions imbedded in observer-based design approaches, and, because it is adaptive, to better emulate the anticipated intercept behavior of biological systems. Adaptation can also be used to introduce other benefits in the engineered guidance and control solution, such as the ability to address system nonlinearities.

3.1 Review of Proportional Navigation

A simple two-dimensional target engagement scenario is depicted in Figure 3.1.




Figure 3.1. A planar engagement model.

The line-of-sight (LOS) rate and range rate for this scenario are given by:

$$\begin{aligned}\omega = \dot{\lambda} &= \frac{V_T \sin(\theta_T - \lambda) - V_M \sin(\theta_m - \lambda)}{R} \\ \dot{R} &= V_T \cos(\theta_T - \lambda) - V_m \cos(\theta_m - \lambda)\end{aligned}\tag{3.1}$$

Defining a_m as the missile acceleration normal to the LOS, then

$$\begin{aligned}V_m \dot{\theta}_m &= a_m / \cos \theta_m \\ V_T \dot{\theta}_T &= a_T / \cos \theta_T\end{aligned}\tag{3.2}$$

To simplify the analysis we introduce the following two assumptions:

Assumption 1: The velocities are constant

Assumption 2: The range rate is constant

Differentiating the first of (3.1), and using $\dot{R} = -\tau \dot{\omega}$, where τ denotes time-to-go, and invoking the above assumptions we obtain

$$\dot{\omega} = 2\omega / \tau + (a_T - a_m) / R \quad (3.3)$$

3.2 Required Acceleration

Substituting the PN guidance law,

$$a_m = -N \dot{R} \omega = N V_c \dot{\omega} \quad (3.4)$$

where V_c is the closing velocity, gives

$$\dot{\omega} = (2 - N)\omega / \tau + a_T / R \quad (3.5)$$

It is well known that PN guidance with navigation gain, $N = 3$, is optimal in the case where $a_T = 0$ (non-maneuvering target)². Optimality is with respect to minimizing the integral square of the acceleration, with intercept formulated as a terminal constraint. The resulting acceleration against a constant accelerating target is given by

$$a_m / a_T = \frac{N}{N-2} \left[1 - \left(1 - t/t_f \right)^{N-2} \right] \quad (3.6)$$

Figure 3.2 illustrates the normalized acceleration resulting from (3.6) as it depends on N and t/t_f .

Augmented PN guidance (APN) is optimal for constant target maneuvers, and is given by

$$a_m = N(\dot{R} \omega + a_T / 2) \quad (3.7)$$

The resulting acceleration against a constant accelerating target is given by

$$a_m / a_T = 0.5N(1 - t/t_f)^{N-2} \quad (3.8)$$

Figure 3.3 illustrates the normalized acceleration for this guidance law. Note that the acceleration profiles are monotonically increasing functions of t/t_f for PN, and the reverse is true for APN. Bear in mind that for $N = 3$ and $a_T = \text{constant}$, the APN law is optimal, whereas the PN law is not.




Figure 3.2. Normalized acceleration due to target maneuver for PN guidance.




Figure 3.3. Normalized acceleration due target maneuver for APN guidance.

3.3 The Effect of Target Maneuvers

The above laws, when implemented in their ideal forms, result in zero miss against most forms of evasive target maneuvers. Real implementations result in non-zero miss distances due to a combination of factors. The primary sources of miss are: lags within the guidance loop, nonlinearities, sensor noise and other sources of tracking error such as radome slope error. Sensor errors also place a lower limit on the effective time constant of the guidance loop³. An illustration of the miss that results when a first order lag with a time constant of $T = 1.0$ sec is modeled in the guidance loop, is given in Figure 3.4. This shows the miss due to a constant $3G$ target acceleration as a function of time-to-go at the start of the maneuver.

While Figure 3.4 illustrates the effect of a constant target maneuver, it can be shown that the optimal evasive maneuver is bang-bang, with reversals occurring at times-to-go corresponding to the peaks in the above curves³. So for $N = 3$, the optimal evasive strategy is to reverse the maneuver at approximately $\tau = 2.5$ seconds. This would have the net effect of doubling the peak miss distance, or achieving a total miss of approximately 54 feet. Kalman filters have traditionally been used in the guidance loop to estimate target acceleration, in order to implement APN. While this is very effective against target maneuvers that belong to the class for which the filter is designed, it also makes the overall design more vulnerable to evasive strategies that take advantage of these modeling assumptions.




Figure 3.4. Miss distance (feet) versus time-to-go at launch due to a $3G$ target maneuver

3.4 Neural Network-Based Adaptive Guidance

There are four (4) major challenges in considering the application of adaptive control to guidance problems. The first concerns the fact that guidance problems are fundamentally finite time problems, whereas adaptive control is formulated as a problem in stability for an infinite time horizon. That is, Lyapunov stability analysis is employed to derive or justify the adaptive laws, and the focus is on stability of an equilibrium point, normally taken as the origin in the error coordinates. One approach (developed in phase I) to overcome this problem is to consider the following time transformation

$$t_1 = \ln(1/\tau) \quad (3.9)$$

where $\tau = t_f - t$. This transforms the LOS rate dynamics in (3.3) to the following linear, time invariant form

$$d\omega/dt_1 = 2\omega + (a_m - a_T)/\dot{R} \quad (3.10)$$

This suggests that the adaptive control problem should be formulated in the t_1 time scale. The adaptive laws would be integrated in the real time scale, but with an integration step of $dt_1 = dt/\tau$, where dt is the increment in real time.

The second difficulty surrounds the fact that the guidance problem is fundamentally an output feedback problem. The adaptive output feedback problem, and its application in design of a high-bandwidth autopilot were reviewed in Chapter 2. In particular, this direct approach to adaptive output feedback design avoids the introduction of a model-based observer⁴. This formulation is also applicable to problems with unmodeled dynamics, however this issue is beyond the scope of the discussion in this chapter. However, a drawback to the method of Chapter 2 and Ref. 3 is that it is based on feedback inversion of the plant dynamics, which constitutes the third major barrier to be surmounted. Feedback inversion is not appropriate for intercept guidance applications for numerous reasons. Ideally, we seek an adaptive approach that augments a nominal intercept guidance solution (such as PN) to correct for evasive target maneuvering. One such architecture, based on fundamental concepts developed under an independent effort at Georgia Tech, which appears to be suitable for surmounting this difficulty is shown in Figure 3.5. This architecture assumes that PN guidance will be augmented by an adaptive process in such a way that the plant behaves like the model response, where the model response is based on Equation (3.10) with $a_T = 0$. In the ideal case, the main differences between these responses are due to the target acceleration, and nonlinearities in the engagement dynamics, part of which may be due to saturation. Therefore, one role of the adaptive process is to augment the PN law by canceling the effect of the target maneuver. We have shown that the error dynamics for this architecture fits the form addressed in Ref 3, and therefore the NN based adaptive control laws developed there should be applicable. *Note that the proposed adaptive guidance architecture is appropriate for augmenting any nominal guidance solution, such as that derived from the study of biological systems.*




Figure 3.5. Adaptive guidance architecture

The fourth technical challenge is to modify the adaptive process to permit adaptation to continue during periods of control saturation. Adaptive processes are highly sensitive to saturation, similar to integrator windup in non-adaptive systems. Ref. 5 describes a method called 'hedging', which can be used to prevent the adaptive process from attempting to cancel selected effects that can be represented at the input side of the plant. These include effects due to saturation, and command limits that are intentionally inserted to avoid exceeding airframe limits. Hedging works by moving the model response backwards by an amount equal to the deficit in achievable acceleration. The location where the hedging is inserted is illustrated in Figure 3.5. In this case, the hedging signal amounts to the difference between the command missile acceleration (a_{mc} in Fig. 3.5) and the actual missile acceleration. This means that all effects due to vehicle + autopilot dynamics, including saturation, are hedged. Thus, the NN is prevented from adapting to these effects. Subsequently, we may wish to adapt to the airframe dynamics, and hedge only the effects of saturation. This will be explored during the Phase II part of this effort.

3.5 Numerical Results

We have conducted a preliminary test using the architecture in Figure 3.5 and the adaptive control approach of Ref. 4. A second order model was used to represent the autopilot lag in the guidance loop, with $\omega_n = 2/T$ and $\zeta = 2/T / \omega_n$. This choice results in T being the effective guidance loop time constant, in the sense that the second order model approximates the rise time of a first order system with this time constant. We chose $T = 0.5$ for the study. For this model, the optimal target maneuver against a PN guidance law with $N = 3$ is a reversal at time-to-go = 1.15 seconds, close to the value predicted by Figure 3.4 for a first order lag if the time axis is viewed as normalized by $1/T$. The PN guidance law alone resulted in a miss distance of

41.6 feet. With adaptive guidance, the miss distance was reduced to 2.8 feet. This result may be further improved if we correctly account for the relative degree of the LOS rate, thereby including the effect of the autopilot lag in the design, rather than simply hedging the acceleration command for these dynamics.

Figure 3.6 compares the resulting acceleration profiles in g's. Note that the adaptive controller appears to be effective in estimating target acceleration. Following an initial transient, its acceleration profile decreases with time, similar in form to the ideal guidance law result that has perfect knowledge of target acceleration (see Figure 3.3). Note that the initial response of the non-adaptive PN law is similar to that depicted in Figure 3.2.

In Fig. 3.7 we examine the effect of an acceleration command limit of 5 g's, for the case of PN guidance ($N = 3$) with an autopilot lag of 0.5 seconds, and a constant 3 g target maneuver. Because the PN law results in an acceleration profile that increases linearly with time, the limit ultimately results in a miss distance of approximately 209 feet. The adaptive law results in negligible miss distance. These results demonstrate the ability to correct for target acceleration without the use of an observer, or pre-defined target acceleration model. The approach also corrects the guidance law for the effect of nonlinear engagement dynamics that result in large rotations of the LOS angle. However, this is not apparent from the results in this example.




Figure 3.6. Comparison of acceleration profiles.




Figure 3.7. Comparison of acceleration profiles for a 5-g limit.




Figure 3.8. Comparison of trajectories for a 5-g limit.

Fig. 3.8 shows a comparison of the trajectories when subjected to a 5 g limit. Note that the trajectory with adaptive guidance does a better job of leading the target, consequently the acceleration needed near the end of the trajectory is significantly reduced, making it far less vulnerable to a reversal in the target maneuver. The miss distances are not apparent at this scale.

Figure 3.9 shows the effect of reducing the missile acceleration limit to 3.5 and 3.0 g's in the adaptive guidance solution. Note that with hedging, adaptive guidance behaves quite well, even when the maneuver advantage of the missile is slightly greater than that of the target. However, when the maneuver capability is matched to the target, then the adaptive solution is not able to compensate the PN law for target acceleration at this initial engagement geometry (head-on). Since, the target can not be intercepted from this condition even with maximum acceleration, it is doubtful that this result can be significantly improved without either accounting directly for the lag in the guidance loop, or investigating the possibility of recovery after a miss has occurred. The degree to which this can be done will be one topic for future investigation.




Figure 3.9. Comparison of trajectories' acceleration profiles when the adaptive law is limited to 3.5 and 3.0 g's.

In summary, a novel method had been developed which treats the problem of target acceleration by incorporating a neural network in place of an observer as an element in a guidance law. This results in a direct adaptive approach for compensating a PN guidance law, which should be less sensitive to smart target evasive maneuvers. Preliminary simulation results indicate that the approach has merit.

4. CHARACTERIZATION OF FLESH FLY PURSUITS

The overall goal in this area is to fully understand how male flesh flies track their targets, especially in environments cluttered with obstacles and alternative targets (other flies), and subsequently to use this understanding to emulate this capability in an engineered system. During Phase I, Prof. Cole Gilbert and Dr. P.O. Zanen of Cornell University developed a reliable system for eliciting aerial pursuits of females by male flesh flies (see Figure 4.1). The current flight arena (see Figure 4.2) has a calibrated active volume of 17x17x40cm, which is a subvolume of a larger cage (50x50x80cm) in which single males and 5-8 virgin females are released several days after eclosion as adults. Flies are 15mm long or shorter with a total mass of 70mg (head: 10.4mg). As the participants fly through the active volume, we capture pursuits with a digitizing precision of 1mm. To increase accuracy in digitized orientation of the sensor (head) we attach a small fiducial stick of balsa wood (8mm, 1.6mg) that indicates the seeker's line-of-sight (LOS). Pursuit sequences are captured with a Kodak MotionCorder SR500 at 250fps and directly read into the hard drive buffer of the SR500. Good sequences are archived onto Hi-8 videotape for offline digitizing and subsequent analysis by custom software [Mantid 32, Synceros Inc.] developed by Dr. P.O. Zanen in the Cornell group. This software enables digitization of any objects in the fields of view.

FLESH FLY

Neobellieria bullata (Sarcophagidae)

body length: \leq 15 mm

mass: 70 mg

linear velocity: \leq 1 m s⁻¹

angular velocity: \leq 1000°s⁻¹




Figure 4.1. Male flesh fly under study.

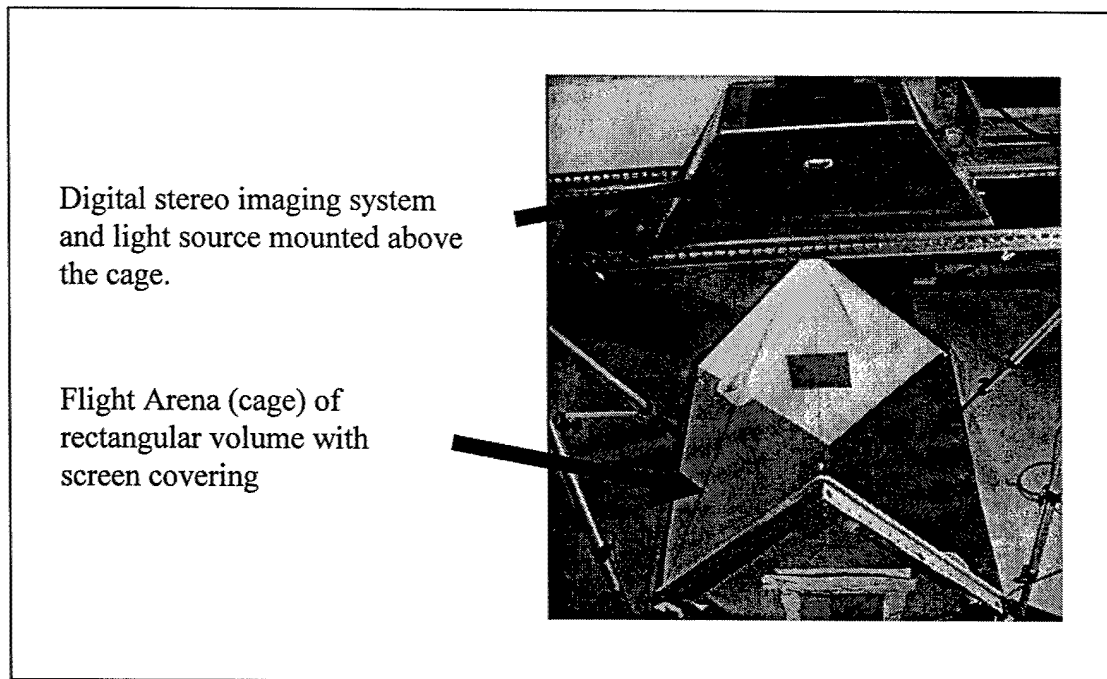


Figure 4.2 Photograph of the current flight cage with imaging equipment mounted above.

During the phase I effort, the test setup and associated procedures described above were gradually improved to ultimately obtain a number of very good pursuit/evasion trajectories that can be analyzed. Figure 4.3 presents a three-dimensional plot of a "straight" trajectory and depicts the path of both the male pursuer (in black), and the female evader (in red). Points on the plot represent position samples at distinct points in time. The male has been fitted with a fiducial stick as described above and, though hard to see at this scale, the figure includes the line of sight angle by means of a vector ascribed to the male position at each sample time. Figure 4.4 presents a second "curved trajectory. Figure 4.5 presents a sequence in which the female successfully evades the male, but in which the male successfully reengages his target to complete the chase. Each of these figures is accompanied by a digital image sequence which is much more effective for initial description of the scenario, but which cannot be included in a printed document such as this. The video clips of these three trajectories were presented to the sponsor at the phase I final program meeting, and delivered on a compact disk. Also included in Figure 4.5 is a plot of the computed position of the female in the field of view of the male during the pursuit. Analysis of the optical flow of this type of data as it evolves during the pursuit is planned for the next phase, and will reveal much about the sensory input to the male during the pursuit.




Figure 4.3 Straight Trajectory




Figure 4.4 Curved trajectory





Figure 4.5 Evasive target with correction by pursuer on left, and plot of computed target position in field of view of pursuing male on right.

The team made significant progress in this area in Phase I, and clearly demonstrated the feasibility of collecting accurate three-dimensional position and orientation data for analysis. However, further improvements to the test procedure are still needed to fully characterize the engagement. The proposed improvements are discussed in detail in Section E of the Phase II proposal, as are the plans for experiments which will reveal the response of the male in pursuit to background clutter and obstacles.

5. IMAGE PROCESSING ALGORITHMS

This section summarizes visual tracking techniques that have been improved and applied as part of the Phase I effort. This work draws upon team member Prof. Allen Tannenbaum's experience in image processing and computer vision employing certain geometric variational evolution equations.

Vision is a key sensor modality in both the natural and man-made domains. The prevalence of biological vision in even very simple organisms, indicates its utility in man-made machines. More practically, cameras are in general rather simple, reliable passive sensing devices which are quite inexpensive per bit of data. Furthermore, vision can capture multispectral information at a high rate at high resolution, and with a wide field of view. Finally, cameras can be used in an active manner. Namely, one can include motorized lenses mounted on mobile platforms that can actively explore the surroundings and suitably adapt their sensing capabilities. For some time now, the role of control theory in vision has been recognized. In particular, the branches of control that deal with system uncertainty, namely adaptive and robust, have been proposed as essential tools in coming to grips with the problems of both biological and machine vision. These problems all become manifest when one attempts to use a visual sensor in an uncertain environment, and to feed back in some manner the information gathered.

Specifically, Dr. Tannenbaum's research has impacted the following areas:

(a) Geometric and Statistical Methods for Deformable Contours:

One of the key techniques in active vision and tracking is that of deformable contours or snakes. These are autonomous processes that employ image coherence in order to track features of interest over time. Our general approach is based on the theory of geodesics and minimal surfaces in a conformal geometry. We have introduced some more sophisticated features into our geometric-based cost function. In particular, we incorporate non-local information into our models based on the use of adaptive filtering schemes and Bayesian statistics. This leads to a natural synthesis of our work in knowledge-based segmentation and active contours.

(b) Area-Preserving Mappings and Optical Flow:

The computation of optical flow has proven to be an important tool for problems arising in active vision, including visual tracking. The optical flow field is defined as the velocity vector field of apparent motion of brightness patterns in a sequence of images. It is assumed that the motion of the brightness patterns is the result of relative motion, large enough to register a change in the spatial distribution of intensities on the images. We have considered various constrained optimization approaches for the purpose of accurately computing optical flow. In phase I, we have weakened the optical flow constraint with one based on the theory of area-preserving mappings. We want to also include an optical flow term in our active contour functional to better track moving images.

(c) Optimal Transport and Image Registration:

Image registration is the process of establishing a common geometric reference frame between two or more data sets from the same or different imaging modalities taken at different times. This is essential for data fusion, since registration can provide automated methods that align multiple data sets with each other and with a given object of interest. We have been using these methods for data fusion as well as a new approach to optical flow for tracking. Our measure of similarity can be based on comparing the mass densities of the images.

We now summarize some of our findings.

Segmentation. One of the key problems in vision is segmentation, i.e., the partitioning of an image into homogeneous regions with "homogeneity" defined in various contexts, e.g., texture, intensity, color, etc. Numerous approaches have been proposed for this. We have considered methods using a synthesis of partial differential equation methods based on the calculus of variations and statistical based ideas.

We have used a combination of anisotropic diffusion and the Bayesian methods for the segmentation of such difficult noisy imagery such as SAR (synthetic aperture radar) and ultrasound. In our approach, we introduce a priori knowledge about the number of objects present in the image via Bayes' rule. Posterior probabilities obtained in this way are then anisotropically smoothed, and the image segmentation is obtained via MAP classifications of the smoothed data.

The model we employ begins with the assumption that the image is composed of n classes of objects. The goal of our segmentation is to determine to which class each pixel in the image belongs. We assume that the value of each pixel in a given class can be thought of as a random variable with a known normal distribution, and that these variables are independent across pixels. Given a set of intensity distributions and priors, we can apply Bayes' Rule from elementary probability theory to calculate the posterior probability that a given pixel belongs to a particular class, given its intensity. We can then calculate the posteriors P , and then to apply anisotropic smoothing to each P . Specifically, we have chosen to smooth by evolving P according to a discretized version of the affine geometric heat equation. This particular diffusion equation was chosen because of its affine invariance, because it preserves edges well, and because of its numerical stability. The final segmentation is obtained using the maximum *a posteriori* (MAP) probability estimate after anisotropic smoothing.

These methods fit very naturally into our overall research paradigm of combining curve evolution with statistical techniques. They will be carefully implemented and tested in phase II in order to develop the optimal combination of the statistical and variational-based segmentation approaches. Furthermore, temporal information may be easily included in this scheme as is shown by the images below (Figure 5.1). These images were generated in the phase I effort by processing computer generated imagery from a real-time missile/target simulation as is proposed

for the hardware-in-the-loop testing in Phase II. Clearly identified and tracked in the image are the horizon and the target (both indicated in red).




Figure 5.1 Two frames of automatic tracking of a fixed-wing UAV target in flight using active contours with statistical methods.

Optical Flow. The computation of optical flow is an important tool for problems arising in active vision. The optical flow field is the velocity vector field of apparent motion of brightness patterns in a sequence of images. One assumes that the motion of the brightness patterns is the result of relative motion, large enough to register a change in the spatial distribution of intensities on the images. Thus, relative motion between an object and a camera can give rise to optical flow. Similarly, relative motion among objects in a scene being imaged by a static camera can give rise to optical flow.

A typical method is to consider a spatiotemporal differentiation method for optical flow. Even though in such an approach, the optical flow typically estimates only the isobrightness contours, it has been observed that if the motion gives rise to sufficiently large intensity gradients in the images, then the optical flow field can be used as an approximation to the real velocity field and the computed optical flow can be used reliably in the solutions of a large number of problems. Thus, optical flow computations have been used quite successfully in problems of three-dimensional object reconstruction, and in three-dimensional scene analysis for computing information such as depth and surface orientation. In object tracking and robot navigation, optical flow has been used to track targets of interest. Discontinuities in optical flow have proved an important tool in approaching the problem of image segmentation.

One constraint which has often been used in the literature is the "optical flow constraint" (OFC). The OFC is a result of the simplifying assumption of constancy of the intensity, $E =$

$E(x,y,t)$, at any point in the image. It can be expressed as the following linear equation in the unknown variables u and v :

$$E_t + E_x u + E_y v = 0.$$

Here u and v are the x and y velocity components of the apparent motion of brightness patterns in the images, respectively. It has been shown that the OFC holds provided the scene has Lambertian surfaces and is illuminated by either a uniform or an isotropic light source, the 3-D motion is translational, the optical system is calibrated and the patterns in the scene are locally rigid.

It is not difficult to see that computation of optical flow is unique only up to computation of the flow along the intensity gradient of E . This is the celebrated *aperture problem*. One way of treating the aperture problem is through the use of regularization in computation of optical flow, and consequently the choice of an appropriate constraint. A natural choice for such a constraint is the imposition of some measure of consistency on the flow vectors situated close to one another on the image.

The optical flow constraint above is of course very strong. Motivated by work on area-preserving maps, we have considered in phase I a modification that could be placed in a variational setting. Indeed, given a family of nowhere-zero 2-forms τ' we have an explicit method to determine a family of diffeomorphisms ϕ' such that

$$(\phi')^* \tau' = \tau^0.$$

Differentiating this expression yields an expression very similar in form to the standard optical flow constraint with a divergence term added. Optical flow computed with respect to this new expression seems to give rather encouraging results as is illustrated by Figure 5.2. These are data frames generated in Prof. Gilbert's lab during Phase I and are images of a male flesh fly in pursuit of a female within a caged area and including another female. The red vectors represent the optical flow field associated with the motion of the three flies.

Optimal Transport, Image Registration, and Optical Flow. In this section, we propose a method for image registration and warping based on the classical problem of optimal mass transport. The mass transport problem was first formulated by Gaspar Monge in 1781, and concerns the optimal way (in the sense of minimal transportation cost) of moving a pile of soil from one site to another. This problem was given a modern formulation in the work of Kantorovich, and so is now known as the *Monge-Kantorovich problem*.

This type of problem has appeared in econometrics, fluid dynamics, automatic control, transportation, statistical physics, expert systems, and meteorology. It also naturally fits into certain problems in computer vision. In particular, for the general tracking problem, a robust and

reliable object and shape recognition system is of major importance. A key method to carry this out is via *template matching*, that is one wants to match some object to a given catalogue of objects. Typically, the match will not be exact and hence some criterion is necessary to measure the "goodness of fit." The matching criterion can also be considered a *shape metric* for deciding the similarity of one shape to another.




Figure 5.2 Two frames of automatic tracking of flies using optical flow.

We are interested in applying these techniques to the problem of registration. Indeed, many registration problems can be formulated in terms of mapping densities. The key idea is to find the optimal mapping via an equivalent problem involving certain factorizations (called "polar") of mass-preserving mappings. We have done this via a natural gradient descent technique similar to the methods that we considered for area-preserving maps of minimal distortion.


It is important to note that there are a number of ways in which multimodal registration impacts image-guided tracking and control systems. First, it allows for fusing the different types of information of each imaging modality, providing better and more accurate information than each image viewed separately. Further, it allows quantitative comparison of images taken at different times, from which precise information about evolution over time can be inferred. Third, it allows for updating in real time a pre-existing image or model with dynamic data.

Multimodal registration is in fact not one but many problems. It greatly depends on the imaging modalities to be matched, on the objects being imaged, and the requirements of the given task. These parameters determine the assumptions and technical requirements, which in turn greatly influence the choice of models and solution methods.


Multimodal registration proceeds in several steps. First, each image or data set to be matched should be individually calibrated, corrected for imaging distortions and artifacts, and cleared of noise. Next, a measure of similarity between the data sets must be established, so that one can quantify how close one image is from another after transformations are applied. Such a measure may include the similarity between pixel intensity values, as well as the proximity of predefined image features such as fiducials, landmarks, surface contours, and ridge lines. Next, the transformation that maximizes the similarity between the transformed images is found. Often this transformation is given as the solution of an optimization problem where the transformations to be considered are constrained to be of a predetermined class. Once the transformation is obtained, it can be used to fuse the image data sets in order to plan a task, to navigate, or to track.

In our case, we are comparing densities on images via the Kantorovich-Wasserstein distance. This occurs, e.g., in functional imaging where one wants to compare mass densities of various features deforming over time, and obtain the corresponding elastic warp map. It also allows us to compare given scalar and vector fields on images.


There have been a number of algorithms considered for computing an optimal transport map. For example, methods have been proposed based on linear programming (for measures which are weighted sums of delta functions) and on Lagrangian mechanics closely related to ideas from fluid flows (Euler equation, etc.) In our work, we have used a gradient descent approach based on the notion of *polar factorization* which allows one to remove the curl from a vector field in computing the optimal transport map. These ideas can naturally be used in studying optical flow as is illustrated in the deforming heart image given below.



(a)



(b)



(c)

Figure C.4 Optimal transport based optical flow. (a) is original image with flow lines indicate. (b) is the target image. (c) shows the warping map.

6. SUMMARY, CONCLUSIONS, AND RECOMMENDATIONS

The phase I effort documented herein produced an innovative and successful method for direct adaptive guidance for intercept of a maneuvering target without the use of a model of target acceleration. Furthermore, the formulation allows for neural network augmentation of a variety of nominal guidance strategies, such as those that arise from our study of biological systems. Feasibility of the new adaptive guidance method was clearly demonstrated in simulation of a 2-D target engagement.

The phase I effort also produced a very successful experimental method for capturing the trajectories and orientations of flies in pursuit of a target, and clearly demonstrated the feasibility and potential of collecting the 3-D and orientation data necessary to fully characterize the tracking, guidance, and control strategy employed by the flesh fly.

The phase I effort also produced improvements to, and demonstration of, practical image processing algorithms that can be used to realize the envisioned system for controlled active vision in hardware and software. The algorithms were exercised using simulated imagery of a 2-D missile engagement with proportional navigation against a maneuvering target, and was successful in autonomously selecting and tracking the target. Live demonstrations of this tracking algorithm were presented to the sponsor at the last program review meeting. The computation of optical flow was also improved, and applied to digital image sequences of the male flies in pursuit. The tracking algorithm was shown to be very effective in this application despite the cluttered environment and obstacles that temporarily obscured the flies from view. In particular, the tracking algorithm's ability to recognize and recover tracking after the fly had been temporarily obscured from view was demonstrated.

It is recommended that these diverse technical elements be combined in follow-on research to produce, integrate, test and demonstrate practical biologically-inspired methods for autonomous guidance and control that engage targets in densely-cluttered environments using vision as the primary sensor modality.

REFERENCES

1. Calise, A., Hovakimyan, N., and Lee, H., "Adaptive Output Feedback control of Nonlinear Systems Using Neural Networks," American Control Conference, 2000.
2. Bryson, A.E., and Ho, Y.C., *Applied Optimal Control*, Blaisdell, Waltham, MA, 1969.
3. Zarchan, P., *Tactical and Strategic Missile Guidance*, Progress in Astronautics and Aeronautics, A. Richard Seebass, Volume 124, AIAA, Washington, DC, 1990.
4. Calise, A. J., N. Hovakimyan, and M. Idan, "Adaptive Output Feedback Control of Nonlinear Systems Using Neural Networks," Submitted for Publication in *Automatica*, January 2001.
5. Johnson, Eric N., A.J. Calise, and H. El-Shirbiny, "Feedback Linearization with Neural Network Augmentation Applied to X-33 Attitude Control," AIAA Guidance, Navigation, and Control Conference, Denver, CO, 2000.